High Speed Token Ring Alliance

September 23rd - 25th, 1997
Raleigh/Durham, NC

Minutes of Technical Meeting #1

Neil Jarvis
Microvitec plc, R'n'D, York, UK
Neil.Jarvis@mail.york.proteon.com
http://www.york.proteon.com/hstr/

Revision History		
v1.0	29 Sep 97	First Release

Contents

Attendees

Karl Reinke, Bay Networks
Paul Gessert, Bay Networks
Ken Wilson, IBM
Tam Ross, IBM
Bob Love, IBM
William Ellington, IBM
Jeffrey Lynch, IBM
Kathy Wilhelm, IBM
Richard Knight, Madge
Dave Wilson, Madge
Andy Fierman, Madge
Earl Crandall, Micro Linear
Edgardo Laber, Micro Linear
Neil Jarvis, Microvitec
Shay Agmon, Novacom
Stan Kmiec, SilCom
John Stephen, SilCom
Ivoa Jeppeson, Olicom
Ole Hansen, Olicom
Benny Jensen, Olicom
Bo Thomsen, Olicom
Keith Luke, Pulse
Brian Buckmeier, Pulse
George Duanne, 3Com
Mark French, Treseq
Tedd Frechette, Cabletron Systems
Carson Stuart, Cisco Systems
Rafael Mantilla Montalvo, Cisco Systems
Michael See, Xylan

Agenda

Tuesday	
9:00 9:15	Introductions
9:15 5:00	What to standardise? IBM Presentation on specific changes to DTR standard for MAC and changes to 100 BaseT sections for PHY portion of draft.
Wednesday	
8:00 5:00	Presentation by Bay Networks and possibly others on potential implementations Continued in-depth discussion on proposed first draft of new standard
Thursday	
8:00 4:00	Paper by Cisco (802.1 considerations for 802.5) Concluding discussions on developing proposed ??

Meeting Objectives

Technical discussion only. This means no discussion of PARs etc.

Decide what it is we want to standardise. There are already two interpretations of which area should be standardised; at the signal level (i.e. what symbols are seen on the media), or at the MII interface level. These are sufficiently different that we need to decide which area we want to talk about.

Ken gave a quick recap on his August presentation. Key points are full duplex only, 4/16/100 switch/adapter configurations. Quick implementations, using as much existing silicon as possible. CAT5/STP only. No CAT3. PAR says use TX PHY devices.

Differing camps:
1) Standardise symbols on 100BaseTX for token ring. This implies an MII interface.
2) Standardise 100Mbit/s MII
3) Standardise 4/16/100Mbit/s MII

There is one common area for all these camps; the symbols on 100BaseTX for token ring.

Presentations

Ken Wilson, IBM

Ken gave a quick recap of the changes to 802.5 he proposed during the August 802.5 interim meeting.

Edgardo Laber, Micro Linear

Edgardo showed how the current definitions of the MII control and status registers, can be re-used for high speed token ring.

MII Control register definitions

Bits	*Default*	*Enable for force speed*
0.15	0	0
0.14	0	0
0.13	1	1
0.12	1: with AN 0: no AN	0
0.11	0	0
0.10	1: Normal 0: Repeater	0
0.9	0	0
0.8	1	1
0.7	0	0
0.6:0.0	x	x

Auto Neg.	Speed	Duplex	Meaning
0.12	0.13	0.8	
1	x	x	Auto negotiate
0	0	0	10Mbit/s, half duplex
0	0	1	10Mbit/s, full duplex
0	1	0	100Mbit/s, half duplex
0	1	1	100Mbit/s, full duplex

It was agreed that the editors should create a new option flag to enable/disable force speed. Only enable will be allowed in the first draft of the standard. The state of the flag determines which PHY configuration primitives are generated during set_initial_conditions.

Tam Ross, IBM

RX_ER within a frame are counted once (with line_error). RX_ER outside a frame are counted once with a new counter (idle_error ?).

How to we reference 802.3? Bob Love to investigate.

802.5 selector field for auto-negotiation is to be requested by Bob Love.

Novacom

Defines an implementation, where 100Mbit/s MII shares pin assignments with a 4/16 PHY interface. This is a valid implementation, but beyond the scope of the standard. The actual MII signal definition for Token Ring will be defined by the new standard, and were extensively discussed during this meeting.

Karl Reinke, Bay Networks

- Standardise connector pinouts to be the same for 4/16 and 100 (e.g. RJ45 3/6 and 4/5).

- Keep 18200 max frame size.
- No phantom at 100.
- Get 802.5 MII selector value
- Use MLT-3
- Support “Reduced-MII”
- Watch out for 18 byte vs. 64 bytes minimum frame size

Neil Jarvis, Microvitec

Presented a design for a software implementation (invented with Dave Wilson, Madge) of phantom for 100Mbit/s TXI token ring.

Presented a compromise where both hardware and software phantom implementations could co-exist.

Cisco

Presentation about 802.1Q. Deferred until November IEEE 802.5 meeting.

High Speed Token Ring Issues

1) Support auto-negotiation vs. Force speed only

- AN/Parallel detect will not work into a hub. The data stream is repeated back to you, you will open at your capabilities. This will fail if you are a 4/6/100 station into a classic 4/16 hub.

Standardise force speed. We may want to define auto-negotiation for token-ring, but reserve for future standarisation all the bits. Get a 802.5 selector valued assigned by 802.3.
[Resolved]

2) Fix reference to 802.3 PHY by date & version vs. Track future changes

Reference by date and version. Bob Love to confirm with editors.
[Resolved]

3) Use link_status vs. Timer to start registration

Use link_status with a failure timer. TSIS remains.
[Resolved]

4) Standardise rate auto-detect vs. Ignore issue in standard.

Ignore issue in standard.
[Resolved]

5) Wire fault. Do we need it?

Go away and decide if this is a required function. This all depends on whether phantom is implemented in hardware or software.
<Ken to remove wire fault>
[Unresolved]

6) 4/16 PHY interface to MII

Reviewed Novacom presentation.
Outside the scope of the standard.
[Resolved]

7) Lobe Media Test/Repeat Path

There are 2 proposals:

i) Frame repeat path:
   How does the Port repeat a bad frame?
   How long does the Port take to repeat a frame?
   When does loopback happen?
   Is internal test required, since an extended loopback of lobe test now includes the Port?
   Elastic buffer in repeat path.
ii) Error repeat path:
   An above-PHY repeat path could be implemented, with the following characteristics:

   Only repeats error in idles, *not* frames.
   Lobe test needs to be defined in more detail than current standard.
   Does not perform the same style of lobe test as current standard.
   Requires new Annex P.

Frame Repeat Path: Yes 1, No 8, Abstain 10+

Lobe media test will be defined as using error repeat path.

Edgardo will write a paper on new lobe test, and post to reflector.
[Resolved]

8) Phantom

Questions:
How much does Phantom cost to implement?
How much does MAC change if phantom is missing?
What does the customer lose without phantom?
Token ring phantom circuitry is incompatible with ethernet line termination?

Ken and Tam have a suggestion of how to "mimic" phantom. By polling link_status, the port could mimic phantom detect. The station could generate a not_phantom_detect event in the port, by resetting the station PHY. This causes the port to see link_status = 0.

Uses of phantom detection with TXI access protocol by Port:
1) Insert completion indication. [Can be done without phantom]
2) Station remove. [Can be done with a frame]
3) Station power off. [Poll link_status for deassertion]
4) Repeat path request, during recovery. [Repeat lobe test frames in MAC, during lobe test]

Costing for and against phantom in100Mbit/s only implementations

With Phantom	*Without Phantom*
$3.27 per port	A MAC code testing
Needs testing	Needs testing
Shared migration not defined	Migration to shared a concern
Time to market hit	Issue is that 4/16/100 is what is needed
New magnetics for 100 PHY	It is new
	Support fibre/copper and 1Gbit/s

3	13
7 Abstain	

A new approach to the issues:

Port requirements:

100 Mbit/s only allowed	14	
100 Mbit/s only not allowed	3	[6 Abstain]

How would 100Mbit/s shared media work?

- Could be made to work without phantom
- Complexity into concentrator

<Ken will write the draft standard without phantom.>
[Unresolved]

9) Abort Delimiter (see 10)

A number of proposals:

i) ET is a word, containing AI and E, and shared flags (A, C) which could be marked as ‘r’.

ii) Use TX_ER prior to deasserting TX_EN for two nybbles to indicate an abort. Then ET becomes an octet, ‘rrrrrrrr’ for TXI, and ‘ACrrACrr’ for TKP.

- More reliable
- Saves bandwidth

Use proposal ii), TX_ER for two nybbles.
[Resolved]

10) Framing (IFG, size, preamble)

Framing for a good frame transmission:

```
MAC             AC . . . INF0 FCS [with A C & E indications]
TXD<3:0>    xx AC . . . INF0 FCS ET
TX_EN 000000111111111111111111111111110000000000
TX_ER xxxxxx0000000000000000000000000xxxxxxxxx
Wire       I JK data                  TR I
RX_DV 000000111111111111111111111111110000000000
RX_ER aaaaaaxxbbbbbbbbbbbbbbbbbbccxxaaaaaaa
RXD<3:0>    xx AC . . . INF0 FCS ET
```

Framing for an aborted frame transmission:

```
MAC             AC . . . <abort>
TXD<3:0>    xx AC . . . xx
TX_EN 000000111111111111111100000000000
TX_ER xxxxxx0000000000000011xxxxxxxxx
Wire       I JK data     HH TR I
RX_DV 000000111111111111111100000000000
RX_ER aaaaaaxxbbbbbbbbbbccxxaaaaaaa
RXD<3:0>    xx AC . . . xx
```

RX_ER phases:

a - Out of frame, code violation (burst error)
b - In frame, code violation (line error)
c - In frame, abort sequence
x - Won’t happen.

Implementations may require a certain number of idles (inter-frame gap) within a timer. No idles, and the PHY resynchronises. This may limit the frame length. **[Unresolved]**

Use 12 idle bytes after the frame. **[Resolved]**
Use 1 idle byte before frame. **[Unresolved]**

Shared HSTR may require more than 12 idles for an inter-frame gap, to help with the elastic buffer that would be required. **[Beyond scope of PAR]**

There is a concern that since the minimum token ring frame format is shorter then ethernet, an implementation that uses an ethernet 100Mbit/s switch fabric would receive ~2.7 more short frames in burst. Some fabrics may not be able to support this. Implementation dependency. **[Resolved]**

PHY Vendors questions:
1) Length of frame (do long frame upset synchronisation?), MLT-3: more baseline wander with longer frames.
2) Inter-frame gap
3) Preamble requirements
4) Baseline wander
5) Phantom into ethernet connection
6) 18 byte frames OK?

<Tam to present questions to vendors>
<Neil to document framing as described above>

11) Connector pinout

Support 802.3u PHY, except that the connectors shall be
- UTP: RJ45 should be 3/6 4/5, as token ring 4/16.
- STP: As currently defined.

[Resolved]

12) MII nybble order, bit order

Modified 802.3u diagram:

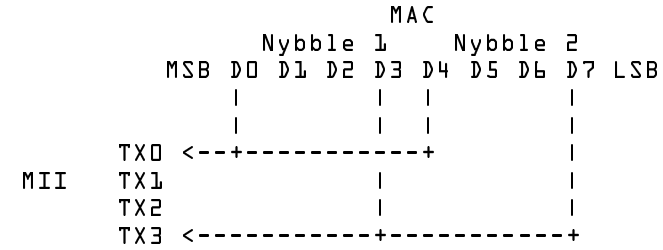

[Resolved]

13) Removing insertion frames

HSTR shared issue only.
[Resolved]

14) 100Mbit/s only Port/NICs. Is this allowed?

Deferred to an 802.5 meeting.
[Deferred]

15) Frame repeat path, and its elastic buffer.

This also requires TKP PH definition. Beyond the scope of this PAR. Vendors are invited to submit proposals on if and how 100 Mbit/s token ring is implemented.

<Ken to ignore this>
[Unresolved]

16) Signal loss

Remove signal loss from TXI.
[Resolved]

17) TKP Access Protocol

Vendors to provide input.
[Unresolved]

Tutorial at next Plenary

Bob is looking for volunteers to help write/present the Tuesday evening tutorial at the Montreal Plenary meeting.

Technical E-Mail Reflector

Dave Wilson has set an e-mail reflector at Madge, to allow technical discussions to take place away from prying eyes.

Address: X-HSTR@dev.madge.com

Next Technical Meeting

We may need another technical meeting prior to the November plenary. Madge, UK have offered to host this meeting. The decision to hold the meeting, will be made after the draft is published and some technical discussion has taken place on the e-mail reflector.

Meeting adjourned